

Frequency Domain Multiplexing for Large-Scale Bolometer Arrays

Helmuth Spieler

Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720

Abstract

The development of planar fabrication techniques for superconducting transition-edge sensors has brought large-scale arrays of 1000 pixels or more to the realm of practicality. This raises the problem of reading out a large number of sensors with a tractable number of connections. A possible solution is frequency-domain multiplexing. I summarize basic principles, present various circuit topologies, and discuss design trade-offs, noise performance, cross-talk and dynamic range. The design of a practical device and its readout system is described with a discussion of fabrication issues, practical limits and future prospects.

Introduction

The next generation of CMB experiments requires a major step up in sensitivity. Wide-field cluster surveys using the Sunyaev-Zel'dovich effect and probes of the E and B-mode polarization anisotropy will become possible through large arrays of several hundred to several thousand bolometers. The development of voltage-biased superconducting transition-edge sensors with strong electro-thermal feedback is a key ingredient that provides well-controlled operating points and stable responsivity over large arrays.¹ A second key ingredient is the development of fully lithographed fabrication techniques.² The remaining challenge is the readout of large arrays, while limiting the number of wires to the cold stage with their attendant heat leaks. One possible solution is time-domain multiplexing.³ Alternatively, one can multiplex in the frequency domain.⁴ That is the topic of this paper.

Basic Principles

Sensors are operated with an AC bias, with a different bias frequency applied to each sensor. When the sensor absorbs signal power, its resistance changes and modulates the signal current. This amplitude modulation transfers the signal spectrum to sidebands above and below the bias frequency. Since each sensor is operated at a different bias frequency, the signals corresponding to different sensors are uniquely distributed in frequency and can be combined in a common readout line. In the off-detector electronics the individual sensor signals are retrieved by a bank of frequency-selective demodulators.

Summing Circuits

The sensor signals can be combined in either a voltage or current summing configuration. First we discuss the voltage summing scheme, as originally published by Yoon et. al.⁴ and shown in Fig. 1. The sensors are coupled to the summing loop by transformers, whose secondaries are connected in series and fed to an output amplifier.

To maintain constant voltage bias, the impedance presented by the transformer primary must be much smaller than the sensor resistance, so $\omega L_p \ll R$. The summing loop also provides a path for cross-talk. Any current flowing in the secondary loop will induce a voltage on the sensor side of each transformer, so the sensor drive current in a given loop will also include components from the other sensor bias generators. The attenuation of a given carrier coupled into any other sensor is

$$A = \frac{1}{n + \frac{Z_i}{\omega L_s}} \frac{\omega L_p}{R},$$

so the requirement for constant bias voltage $\omega L_p \ll R$ also reduces the cross-talk. Isolation beyond this level for a given n requires an increase in the amplifier's input impedance, usually achieved by feedback.




Figure 1: Voltage summing multiplexer

In addition to the frequency-correlated signal current, each sensor contributes a wideband Johnson noise current. As a result, superimposed on the sideband spectrum associated with each bias frequency, the Johnson noise from all other sensors will also appear. To avoid excessive noise buildup one can introduce frequency selective circuits into each sensor loop, as shown in Fig. 2. By appropriate choice of the circuit Q and the spacing between bias frequencies, the Johnson noise buildup can be kept at arbitrarily low values.




Figure 2: Tuned circuits in each sensor limit the loop bandwidth of the sensor's Johnson noise.

Although the circuits shown in Figs. 1 and 2 reduce the number of readout channels, they still require at least one wire per sensor. Introduction of the tuned circuits allows all bias frequencies to be introduced as a “comb” on a common bias line, as shown in Fig. 3. Now the tuned circuits perform double duty; they “steer” the bias drive to the appropriate sensor loop and also limit the bandwidth of the Johnson noise. Fig. 3 also shows a SQUID readout amplifier, employing series feedback to obtain a high input impedance.

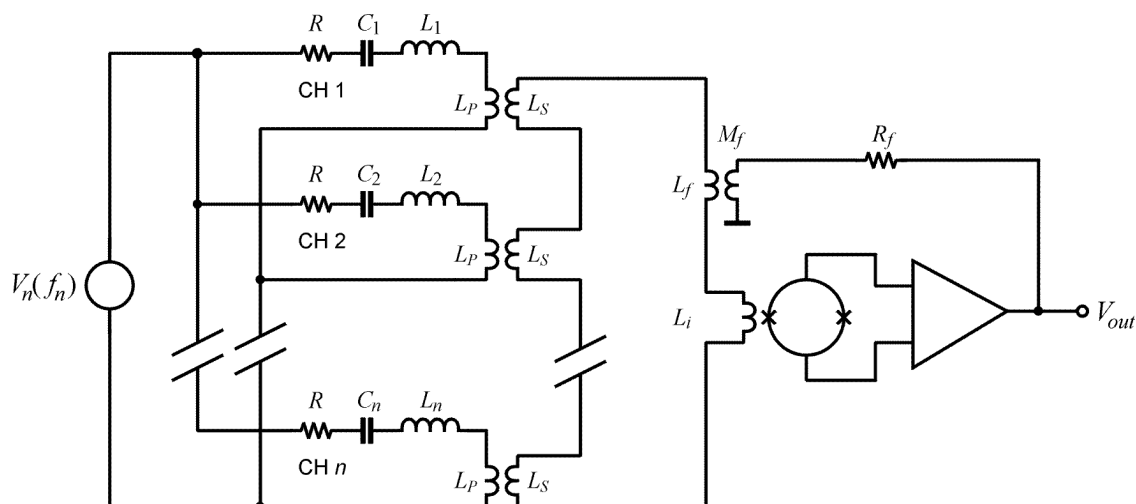


Figure 3: Voltage summing multiplexer with common bias line.

In the current summing scheme (Fig. 4), the individual sensor circuits are joined in a common current summing node at the input of the readout amplifier. To maintain constant voltage on the sensor, the input impedance of the amplifier must be much smaller than the sensor resistance, which is also required for efficient current summing. Two connections read out multiple sensors (enclosed by the dotted box).

In both voltage and current summing circuits that use a common bias line and tuned circuits to “steer” the bias frequencies to the corresponding sensor, the bias-cross talk is determined only by the selectivity of the tuned circuits. The input impedance of the amplifier only introduces an additional impedance into the sensor loop.

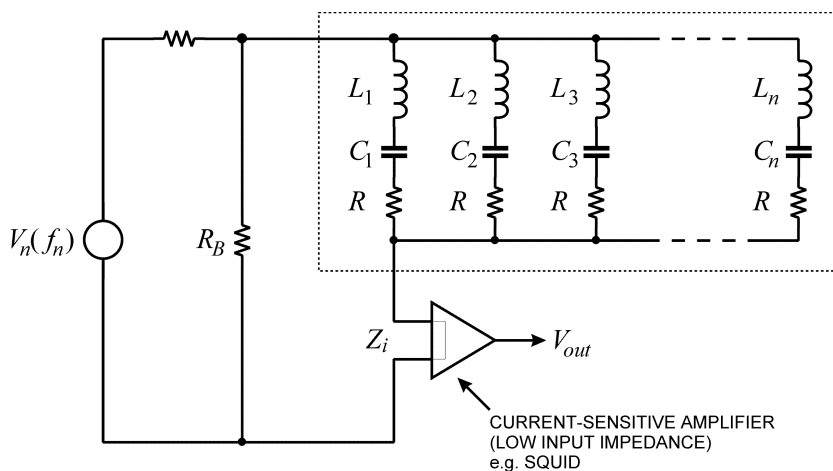


Figure 4: Current summing multiplexer

Apart from the complementary requirements on amplifier input impedance, the voltage and current summing schemes are fully equivalent with respect to electronic noise, dynamic range, and all other properties. The choice between the two is determined by practical considerations.

Circuit Parameters

Cross Coupling of Signal and Noise

In Figs. 3 and 4 the selectivity of the individual tuned circuits determines the cross-talk between channels. For a given bias frequency the maximum current flows through the leg whose resonance matches the bias frequency. However, the finite selectivity of the other legs allows a fraction of this carrier's current to flow through them and be modulated by their respective sensors. The frequency dependence of the current

flowing through the LCR circuit $\frac{I}{I_0} = \left(1 + Q^2 \left(\Omega - \frac{1}{\Omega} \right)^2 \right)^{-1/2}$, where Ω is the normalized frequency

$\Omega = \omega / \omega_0$ and $\omega_0 = 1 / \sqrt{LC}$. I_0 is the current at resonance ($\Omega = 1$) and the circuit $Q = \omega_0 L / R$. At resonance the sensor is voltage biased and the bias power $P = I_0 V$. Off-resonance, however, the voltage across the sensor depends on the fractional current I / I_0 given above and the sensor power $P = I^2 R$. The total cross-talk is the sum of the contributions from all other channels at the bias frequency of interest.

Analogously to signal cross-talk, noise from adjacent sensors is coupled into a given channel through the finite selectivity of the tuned circuits. Noise from the neighbor channels can be introduced in two ways. 1. Each sensor generates a wideband noise spectrum, shaped by the selectivity curve of the corresponding tuned circuit. Thus, noise from all sensors appears at a given bias frequency. 2. The fractional carrier current of frequency f_n flowing through the neighbor channels $f_{m \neq n}$ is modulated by the sensor noise and adds to the total noise in the sidebands associated with f_n . The relative magnitude of this contribution is determined by the bias cross-coupling, as discussed above.

Although the attenuation of the wideband noise depends on the circuit selectivity as does the carrier cross-talk, the relative contributions differ. The ratio of cross-talk power to on-resonance power is equal to the square of the cross-talk current to the on-resonance current and the total cross-talk $(P / P_0)_n = \sum_{m \neq n} (I / I_0)_m^2$,

whereas the noise adds in quadrature, so the total noise $(N / N_0)_n = \sqrt{\sum_{m \neq n} (I / I_0)_m^2}$.

Determination of Circuit Parameters

The maximum number of channels per output line is obtained with a constant bandwidth system using the minimum allowable bandwidth, which must be greater than $f_{\max} = 1 / 2\pi t$ to ensure both constant voltage bias and electro-thermal stability, so $b \approx 2 / t$, where t is the sensor's decay time constant. The bandwidth of a simple LC tuned circuit $b = f_0 / Q$, so constant bandwidth implies that the circuit Q is proportional to frequency. Since $Q = \omega_0 L / R_s$, this is automatically fulfilled if L is kept constant. The spacing between individual frequency "bins" is determined by the required neighbor-channel attenuation.

The circuit parameters are interrelated and determined in the following sequence:

1. The maximum Q is determined by the maximum frequency f_{\max} and the required bandwidth b , $Q = f_{\max} / b$.
A lower Q can be chosen at the expense of channel density.
2. The inductance is set by the adopted Q and the total resistance R in the circuit $L = RQ / 2\pi f_{\max}$.
3. The frequency spacing Δf is determined by the allowable cross-talk. All neighbors m will couple into a given channel n , so $(P / P_0)_n = \sum_{m \neq n} (I / I_0)_m^2$ and $(N / N_0)_n = \sqrt{\sum_{m \neq n} (I / I_0)_m^2}$.
4. The minimum frequency is given by the thermal time constant of the sensor and by the maximum available capacitance together with the inductance calculated in step 2, $f_{\min} = 1 / 2\pi \sqrt{LC_{\max}}$.
5. All of this together determines the number of channels per module $(f_{\max} - f_{\min}) / \Delta f$.

System Bussing

As will be shown below, the number of channels that can be read out by one line is limited. In a large array the MUX “module” can be replicated, using the same circuitry and bias frequencies. Alternatively, one can use oscillator bussing that combines all sensors using the same bias frequency, as shown in Fig. 5. This arrangement also alleviates the channel spacing requirements imposed by bias cross-talk, as different frequencies are on separate busses. However, mismatches of the resonant frequencies in adjacent multiplexer modules will affect the uniformity of the bias drive level. Alternatively, one can gang subsets of channels with larger channel spacings to control cross-talk. This comes at the expense of one oscillator per sensor.

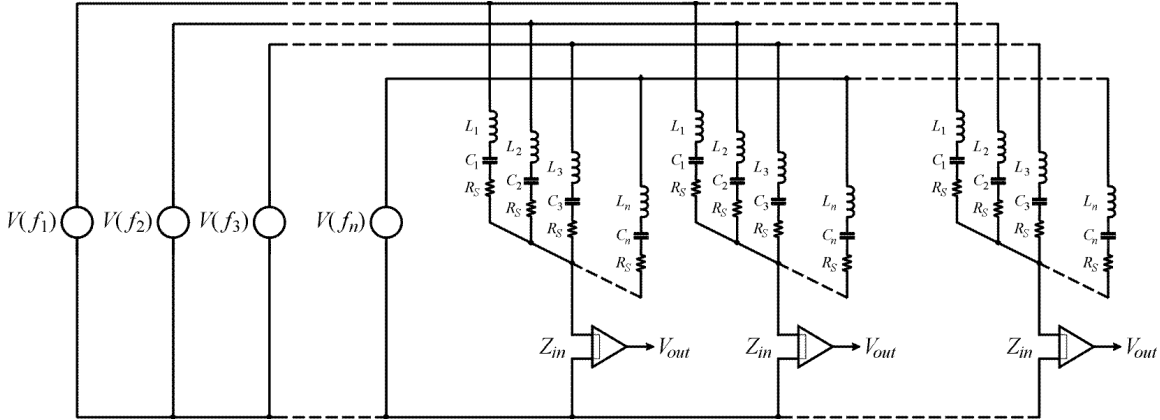


Figure 5: Bussing by common frequency lines to gang multiple multiplexer modules

Assume that the total array has M sensors and that each multiplexer module can combine N sensors into one readout line. Then all configurations require M/N readout lines, but differ in the number of bias lines. Fig. 1 requires M bias lines, Figs. 3 and 4 requires M/N bias lines, and Fig. 5 requires N bias lines. Thus, the total number of lines is $M/N + M$ in Fig. 1, $2M/N$ in Figs. 3 and 4, and $M/N + N$ in Fig. 5. If $M = 1024$ and $N = 32$, the number of lines is 1056, 64, and 64, respectively. Fig. 5 becomes more favorable for large M/N .

Maximum Number of Sensors per Multiplexer Module

In principle, the number of channels in a MUX module can be extended arbitrarily; increasing the number of channels increases the maximum frequency. Since the frequency spacing is constant, high frequencies provide more channel capacity than low frequencies. However, this is limited by two practical constraints, tolerances of component values and the limited dynamic range of the readout.

Mismatch of inductor and capacitor values

In practice, the inductor and capacitor values are subject to fabrication tolerances. Relative tolerances are more critical than absolute deviations. For example, if all capacitor or inductor values are high, all resonant frequencies will shift down and the spacing will remain roughly unaffected. Relative mismatches, on the other hand, can lead to a positive frequency deviation in one channel and a negative deviation in the next higher channel, so that the spacing decreases, with a corresponding increase in cross-talk. Thus, relative component mismatches require increased channel spacings at high frequencies.

Dynamic Range

The dynamic range that can be processed by the amplifier is determined by the noise floor and the maximum signal. Both voltage and current summing multiplexers can provide an ideal noise match from the sensor to the amplifier, so that the overall noise temperature is equal to the sum of the noise temperatures of the SQUID (~ 10 mK) and a single sensor (~ 100 mK). Thus, the noise floor is typically determined by the sensor noise, whereas the maximum signal that can be accepted by the SQUID depends

on the SQUID feedback circuitry. Deliberately mismatching the SQUID input so that the sensor noise barely overrides the SQUID noise will reduce the signal applied to the SQUID and increase the dynamic range. This will be shown below in more detail.

Maximum Acceptable Input Signal

The major problem with SQUID readouts is the limited signal range. Since the net flux excursion in the SQUID is limited to $\pm\Phi_0/4$, feedback systems are critical not just for the stabilization of the operating point, but also to extend the SQUID's dynamic range. SQUID controllers commonly use an integrator to introduce a dominant pole. This does allow the feedback loop to be stabilized against self-oscillation with a single gain adjustment, but at the expense of high-frequency dynamic range. Due to the integrator, the loop gain decreases linearly with frequency over the full frequency range, so the product of maximum input flux and frequency is constant (commonly referred to as a "slew rate limit"). In contrast, a SQUID controller has been designed with constant dynamic range $25\Phi_0$ over the full signal bandwidth of 1 MHz (corresponding to a "slew rate" of $10^8 \Phi_0/s$).⁵ However, a fundamental limit is the finite length of the feedback loop, which even with an ideal single-pole amplifier response limits the loop's gain-bandwidth product.

The SQUID's maximum allowable input signal is increased by the loop gain. For example, a maximum input signal of $25\Phi_0$ requires a loop gain $A_L \approx 25\Phi_0/(\Phi_0/4) = 100$. This loop gain also provides <1% non-linearity for signals $< \Phi_0/6$, which will be adopted as a practical signal limit. To achieve stable operation of the feedback loop the loop gain should roll off to unity at the frequency where the net phase shift is not less than 45° , i.e. one can tolerate an additional phase shift of 135° relative to the 180° phase shift at low frequencies. An ideal single pole system introduces a 90° phase shift above the cutoff frequency. However, because of the finite phase velocity v_p , the physical length l of the feedback loop introduces an additional phase shift $(lf/v_p)2\pi$, so the unity gain frequency is limited to the frequency at which the propagation delay introduces the remaining 45° ($\pi/4$) phase shift, i.e. $f_0 = v_p/8l$. Thus, the gain-bandwidth product of the feedback loop $A_L f_{\max} = v_p/8l$ and the upper cutoff frequency $f_{\max} = v_p/8lA_L$. If $v_p = c$, a conductor length of 10 cm yields a gain-bandwidth product of 375 MHz, so for a loop gain of 100 the upper cutoff frequency is 3.75 MHz. In an insulating medium of dielectric constant ϵ , $v_p = c/\sqrt{\epsilon}$, so for typical coaxial cables $v_p/c = 2/3$ and $l = 10$ cm leads to $f_{\max} = c/8l\sqrt{\epsilon}A_L = 2.5$ MHz. This limit holds for any combination of gain stages in the loop, so it does not improve for series array SQUIDs, for example. To maximize the allowable gain-bandwidth product the propagation delay must be minimized, i.e. by reducing the physical size of the amplifiers and the length of the connecting lines. A localized cold feedback loop is clearly advantageous.

Carrier Nulling

Since the peak signal is dominated by the bias carrier components, one can alleviate the dynamic range requirements by suppressing the carriers. As all of the signal is in the sidebands, suppressing the carrier does not affect signal recovery. This is illustrated in Fig. 6. Suppressing the carriers 10-fold relaxes the loop gain requirements and potentially extends the bandwidth in the above example 10-fold to 25 MHz.

Estimate of Number of Channels per Readout Line

Let the sensor bias and noise currents be I_b and I_n . If the desired noise floor in the SQUID is Φ_n , the mutual inductance of the SQUID input is set to $M_i = \Phi_n/I_n$. Thus, the maximum signal $\Phi_s = n\sqrt{2}I_b M_i$ and the required loop gain $A_L = (\Phi_s/(\Phi_0/6) - 1) \approx 6n\sqrt{2}(I_b/I_n)(\Phi_n/\Phi_0)$. The loop's gain bandwidth product is limited by the length of the feedback loop $f_{\max}A_L \leq c/8l\sqrt{\epsilon}$, so $f_{\max} = c/(48nl\sqrt{2}\epsilon(I_b/I_n)(\Phi_n/\Phi_0))$.

Since the number of channels $n = (f_{\max} - f_{\min})/\Delta f \approx f_{\max}/\Delta f$, one obtains

$$f_{\max} \approx \sqrt{\frac{c}{l\sqrt{\epsilon} 48\sqrt{2}(I_b/I_n)(\Phi_n/\Phi_0)} \Delta f} \quad \text{and} \quad n \approx \sqrt{\frac{c}{l\sqrt{\epsilon} 48\Delta f \sqrt{2}(I_b/I_n)(\Phi_n/\Phi_0)}}.$$

A 10 cm long feedback loop with $\epsilon = 2.25$, $\Delta f = 20$ kHz, $I_b = 1 \mu\text{A}$, $I_n = 10 \text{ pA}/\sqrt{\text{Hz}}$ and

$\Phi_n = 5 \text{ m}\Phi_0/\sqrt{\text{Hz}}$ yields $n = 54$ channels with a maximum frequency of 1.1 MHz. If carrier nulling is used

with a 10-fold suppression, one obtains about 170 channels with a maximum frequency of 3.4 MHz. Note that due to the minimum frequency limit, these channel numbers must be reduced by $f_{\min} / \Delta f$.

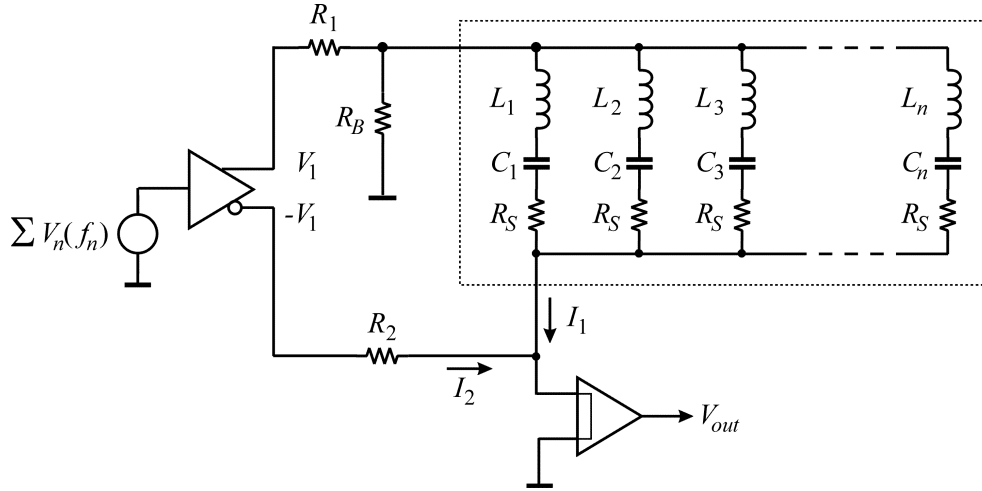


Figure 6: The bias carrier signals can be suppressed to reduce dynamic range requirements.

A Practical Multiplexer Design

A 32-channel multiplexer was designed for use in the circuit topologies shown in Figs. 4 and 5, and fabricated by TRW⁶. The frequency range is 380 kHz to 1 MHz with a nominal channel spacing of 20 kHz. For this test the multiplexer was subdivided into four 8-channel chips, whose resonant frequencies are interleaved so that within one chip the frequency spacing is 4 times the nominal spacing. The 8-channel chips can be used individually or connected in parallel to form a full 32-channel multiplexer. Bond pads are provided to allow measurement of individual inductors and capacitors. Fig. 7 shows the layout of one chip; the die size is 1 cm. The capacitors use Nb₂O₅ dielectric formed by anodization of Nb. This yields a capacitance of 3nF/mm² with very good uniformity. The inductors are implemented as square spirals. Spiral inductors have a large fringing field, so to reduce the sensitivity of the inductance to nearby objects, the inductors utilize a “transformer” geometry as used in SQUIDS, where the inductor spiral is tightly coupled to a square washer slit at one side.⁷ The inductors all have the same value of 40 μH and with a winding pitch of 4 μm are 2.2 mm on a side. The capacitances range from 640 pF to 4.4 nF. At a frequency




Figure 7: Layout of 8-channel MUX chip. The eight inductors surround a central capacitor core.

spacing of 20 kHz and a maximum Q of 500 (at 1000 kHz) the calculated noise degradation due to all neighbor channels is 9%. Channel-to-channel cross-talk is $0.34 \times (Z_{in}/R_s)$, where Z_{in} is the input impedance of the SQUID and R_s is the sensor resistance (0.5 ohms). The factor 0.34 is due to the selectivity of the tuned circuits. For a given sensor it is the sum of the currents at all undesired frequencies divided by the current at the resonance frequency. The devices have been fabricated successfully and are now under test.

Readout Electronics

The frequency-domain multiplexer reverts to full parallelism at the signal recovery stage. In principle, a DSP could be used to retrieve all of the signals, but for the frequencies of interest this is impractical at present. We have adopted a fully analog demodulation system, shown in Fig. 8. The composite signal stream is fed to a bank of phase-synchronous demodulators. The same carrier oscillators used to bias the sensors are used for demodulation. Thus, to first order, oscillator sideband noise cancels out. The local oscillator signal is fed to the mixer through a miniature delay line (not shown) to compensate for the round-trip propagation delay of the cabling between the cryostat and the readout electronics. This maintains the proper phase relationship between the input and demodulation signals. Direct digital synthesis (DDS) oscillators provide both frequency and phase adjustability. The output of the synchronous detector feeds a low-pass filter and then a gated integrator. These integrators are all gated by the same clock, to ensure that all measurements in the array cover the same time intervals. The outputs of the gated integrators are multiplexed to a multichannel PC-based ADC with differential inputs to provide common-mode isolation. Digitization occurs within less time than the integration interval, so there is no dead time. Our current design has 32 channels per multiplex module, so 1000 channels require about 30 modules, all of which are identical. All circuit functions can be performed by economical standard ICs.

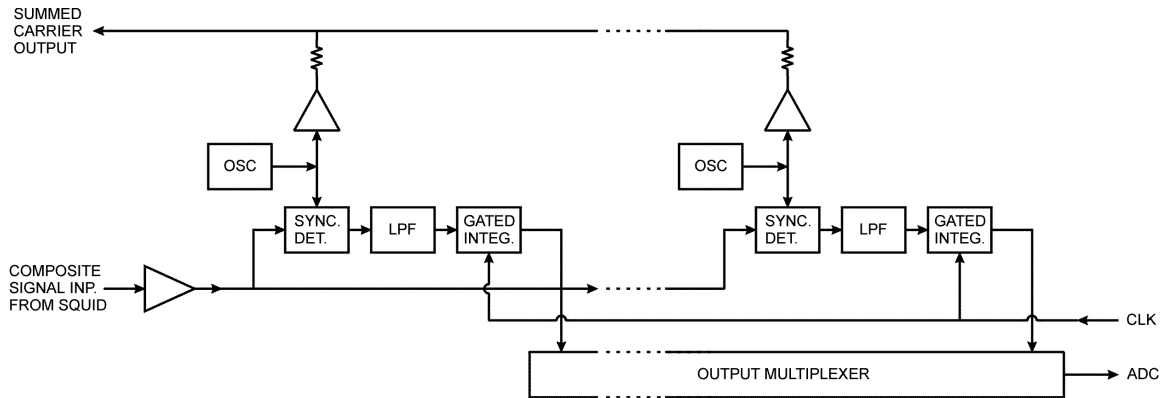


Figure 8: Demodulation and readout electronics

Status and Future Prospects

Although one might read out several hundred sensors with one readout line, this introduces a significant single-point failure mode, so 30 – 50 sensors per readout line appears to be a reasonable choice. Detailed designs are underway for arrays of 300 – 3000 bolometers in next-generation CMB experiments. Use of SQUIDs as readout amplifiers necessitates carefully optimized feedback loops to achieve the required dynamic range and bandwidth. However, SQUIDs are not the only choice. Special MOSFETs designed and fabricated at LBNL operate at 4K and exhibit adequate noise levels at drain currents of $<100 \mu\text{A}$. The signal frequencies used here are well above the $1/f$ noise corner. Use of these devices requires carefully designed matching transformers, but this is greatly facilitated by the low sensor impedance and high-frequency operation. This technology allows a localized feedback loop and extends the large-signal bandwidth limits of a SQUID amplifier. Frequency domain multiplexing has also been demonstrated with x-ray sensors. The circuit in Fig. 2 has been used to multiplex x-ray TES devices with no degradation in resolution.⁸

Summary

Frequency-domain multiplexers provide a practical method of reading out large arrays of superconducting bolometer arrays. They provide continuous readout with no switching transients. Practical designs have been developed and are now being tested and evaluated for use in large-scale CMB experiments.

Acknowledgements: This work was performed within the Berkeley Bolometer Group (P. Richards, A.T. Lee, W. Holzapfel, J. Clarke, *et al.*). A.D. Smith at TRW provided valuable help in the design and fabrication of the multiplexer chips. Support was provided by the Director, Office of Science, Office of High Energy and Nuclear Physics, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

References

1. S.-H. Lee, *et al.*, A Voltage-Biased Superconducting Transition Edge Bolometer with Strong Electro-Thermal Feedback Operated at 370 mK, *Applied Optics* **37**, 3391 (1998).
2. Gildemeister, J.M., Lee, A.T., and Richards, P.L., Monolithic Arrays of Absorber-Coupled Voltage-Biased Superconducting Bolometers, *Appl. Phys. Lett.* **77**, 4040 (2000).
3. Chervenak, J.A., *et al.*, Superconducting Multiplexer for Arrays of Transition Edge Sensors, *Appl. Phys. Lett.* **74**, 4043-4045 (1999).
4. Yoon, J., *et al.*, Single Superconducting Quantum Interference Device Multiplexer for Arrays of Low-Temperature Sensors, *Appl. Phys. Lett.* **78**, 371 (2001)
5. Helmuth Spieler, internal note
6. TRW Inc., One Space Park, Redondo Beach, CA 90278
7. J.M. Jaycox and M.B. Ketchen, Planar Coupling Scheme for Ultra-Low Noise DC SQUIDS, *IEEE Trans. on Magnetics* **MAG-17/1** (1981) 400-403
8. M.F. Cunningham, J.N. Ullom, T. Miyazaki, S.E. Labov, John Clarke, T.M. Lanting, Adrian T. Lee, P.L. Richards, Jongsoo Yoon, and H. Spieler, High-Resolution Operation of Frequency-Multiplexed Transition-Edge Photon Sensors, *Applied Physics Letters* **81** (2002) 159 – 161